

Best Practices in AI and Cognitive Computing

- Marydee Ojala** 2 **Artificial Intelligence Done Right**
Artificial intelligence (AI) has captured the imagination of a wide variety of businesses. I have this image of CEOs in boardrooms around the globe declaring, “We must have AI! Our competitors use AI!...”
- Daniel Fallmann, Mindbreeze** 4 **AI Guidelines for Businesses: Using AI in Your Own Company**
Artificial intelligence (AI) is one—if not the—key technology of our decade. Technological advances in this field are not only fundamentally changing our economies...
- Kelly Koelliker, Verint** 6 **Understand. Anticipate. Improve.**
For decades, organizations have tried to unlock the collective knowledge contained within their people and systems. And the challenge is getting harder, since every year, massive amounts of additional information are created for people to share...
- Scott Parker, Sinequa** 7 **3 Things to Know Before Starting Your AI Journey**
AI-Powered Search Engines—referred to as “Insight Engines” by Gartner and “Cognitive Search” by Forrester—can deliver significant value to organizations these days, provided certain risks are avoided...
- Kamran Khan, Accenture** 8 **From “Searching” to “Finding”: How AI is Unlocking the Power of Unstructured Data**
Unstructured data, which comprises almost 80% of any enterprise’s data, holds untapped value when it comes to addressing challenges and embracing opportunities...
- Mark Gross, Data Conversion Laboratory** 9 **AI and the Building Blocks of Intelligent Content**
The data, information, and analytics economy run on well-curated structured data. No matter your industry—having good curated data and content is critical. It’s increasingly important as more data and content are generated...
- Michael Cizmar, MC+A** 10 **A Best Practice Approach to Insight Engines: 5 Levels of Insight Engine Maturity**
Enterprise search projects start with intentions to provide ‘Google for our organization’ but too often fail to deliver on that promise...
- Chris Caile, Nuance** 11 **IVRs and AI, Unite!**
While interactive voice response systems (IVRs) have been invaluable in reducing contact center costs, we need to be honest: not many are delivering experiences that meet consumer expectations...
- Sean Coleman, BA Insight** ... 12 **Key Considerations in Maximizing the Value of Cognitive Search**
*I am a firm believer in *The 7 Habits of Highly Effective People*, by Stephen Covey. If you’ve not read this book, it is worth the time...*
- Anand Subramaniam, eGain** . . 13 **AI-Powered Customer Service: Use-Cases and Real-World Examples**
Cognitive/AI technologies for customer engagement are white hot. No wonder professionals, who had removed AI from their resumes, are scrambling to add it back in...

Produced by:

**KMWorld Magazine
Specialty Publishing Group**

For information on participating in the next white paper in the “Best Practices” series, contact:

Stephen Faig

Group Sales Director
908.795.3702
sfaig@infotoday.com

LaShawn Fugate

Account Executive
859.361.0667
lashawn@infotoday.com

KMWorld
KMWorld.com

Premium Sponsor



Artificial Intelligence Done Right

By Marydee Ojala, Conference Program Director, Information Today, Inc.



Marydee Ojala

Marydee Ojala is conference program director for Information Today, Inc. She works on conferences such as Enterprise Search & Discovery, which is co-located with KMWorld, and WebSearch University, among others. She is

a frequent speaker at U.S. and international information professional events. In addition, she moderates the popular KMWorld webinar series.

Ojala is based in Indianapolis, Indiana and can be reached at marydee@infotoday.com.

Artificial intelligence (AI) has captured the imagination of a wide variety of businesses. I have this image of CEOs in boardrooms around the globe declaring, “We must have AI! Our competitors use AI! We can’t be left behind!” There might be some table-pounding associated with this scenario. There will certainly be corporate minions scurrying around to fulfill the AI dreams of their CEO.

As Mindbreeze’s CEO Daniel Fallmann assured me in a recent phone conversation, this is not the correct way to go about implementing an AI project. You don’t just go out and buy a box filled with AI that you plug into your computer. There are a lot of moving parts to consider. I asked Fallmann if he recommended starting small rather than trying to take on every possible area of a company where AI could be useful. He corrected me slightly. “I’d say starting focused would be the better way to put it.” In large corporations, for example, Mindbreeze might start with customer service. But it could equally be the maintenance department or research & development.

The idea is to get a 360-degree view of the department’s activities. “Talk with the head of the department,” advised Fallmann, “not with IT.” Look at what information is being collected and the KPIs (Key Performance Indicators) as they are today. What needs to be

where a pharma company’s submission is rejected when a very similar one from a different company is approved.

Fallmann has another concern—dirty data, or as he terms it, “data garbage.” If fed dirty data, a machine learning (ML) application will return lots of false positives, leading to bad decisions. You need 85-90% accuracy, he thinks. Concentrate on the business case and whether the existing data is inaccurate, incomplete, inconsistent, or otherwise compromised, it will be difficult to obtain relevant results.

Mindbreeze Guidelines

In Mindbreeze’s guidelines for businesses, Fallmann advises to start by identifying and determining the use case. Look for opportunities to pinpoint pain points that can be improved through a sustainable AI implementation, but don’t try to tackle too many at once. Then turn to defining the success criteria, which will revolve around business needs, data sources and quality, semantic relationship and extraction, and an ROI calculation. Next is hands-on testing with company data, which may need to be expanded and probably requires some cleansing.

A critical component of the Mindbreeze guidelines is involving the users. In a world defined by continuous change management,

plan for executing on the strategy, making real, measurable, and sustainable improvements that lead to a competitive advantage.

Role of Search

It’s not really a secret that the majority of internal information is in unstructured formats. This has long been a challenge—and a frustration—for those trying to find what they need. Kamran Khan, Managing Director of Search & Content Analytics for Accenture, sees AI as the antidote to that frustration. He points out that AI is actually a “collection of multiple technologies that encompasses deep learning, natural language processing (NLP), search, ML, intelligent sensors, and robotics.” He singles out the use of NLP and ML to understand documents as a key driver to unlock the power of unstructured data.

Although NLP is not a new technology, Khan notes that it’s been going through a transformation recently. The widespread acceptance of web search engines and digital assistants paved the way for NLP applications that understand and respond to natural language queries in a similar fashion. Plus, NLP’s pairing with ML has superseded its earlier dependence on rules written out manually. That increases the scalability and accuracy of document handling.

Khan also thinks that the analysis of unstructured data is becoming more proactive and less reliant on reactive searches. It’s easier to extract relevant information from unstructured data than it was in the pre-AI, NLP, and ML days, which increases the ability of companies to “unlock the full potential of unstructured data.”

Implementing AI and Cognitive

What about AI-powered search engines, aka insight engines or cognitive search? Scott Parker, Senior Product Marketing Manager, Sinequa, has some thoughts. To do AI right, start with a valuable use case. Technology plays a role, but solving a problem takes precedence. Once you’ve got a use case, think about implementing the right methodology by acquiring

changed to make that information actionable? How can AI bring down the workload of the call center or the streamlining of no-defect parts maintenance? In Fallmann’s opinion, it’s not just about saving money, but also about creating a “Wow effect.”

Measuring the result of an AI project isn’t simply an ROI (Return on Investment) exercise. It also needs success criteria that measure the ultimate impact of business process transformation. He gave the example of the U.S. Food & Drug Agency (FDA) and its usage of Mindbreeze to identify inconsistencies in the submissions from pharmaceutical companies. The FDA wants to avoid a situation

feedback from employees will vastly increase your possibilities for success. After all, they know how they do their jobs better than an outsider and can be instrumental in training the data set. A final step in the guidelines is validating the ROI calculation. Moving from a test phase to a departmental project and then to the entire company gives you the chance to measure the improvement in efficiency, employee satisfaction, and overall profitability.

In the final analysis, Fallmann told me that, to do AI right, “you need execution, not just strategy.” When top management calls for an AI strategy, Mindbreeze wants to have a concrete

“The ultimate goal of AI is to get machines that can understand language as humans do. That includes disambiguating words and concepts in a contextual fashion. Understanding what is actually meant when someone says or writes something isn’t as straightforward as you might think.”

data, building a model, deploying and validating your solution, and going into active learning and tuning mode. Parker stresses the importance of choosing a proven and unified technology platform so that you don’t sacrifice context or quality. Assuming you want to achieve your predicted ROI, follow his steps to move from having an AI project as an experiment to a fully implemented, successful initiative.

The dream of uncovering and optimizing the collective knowledge within an organization seems a bit closer now, thanks to AI and cognitive computing, according to Kelly Koelliker, Director of Content Marketing at Verint. The ultimate goal of AI is to get machines that can understand language as humans do. That includes disambiguating words and concepts in a contextual fashion. Understanding what is actually meant when someone says or writes something isn’t as straightforward as you might think. Humans can generally grasp the nuances—although that’s not always true—but computers have a much harder time of it. Advances in the technology, such as semantic and cognitive analysis, lead to better understanding of questions. In turn, that leads to better answers.

Koelliker also points to cognitive computing’s ability to predict what will be asked and to anticipate answers. Enhancing content automatically to incorporate implicit and explicit knowledge improvements has the potential to replace activities such as building new taxonomies or creating specialized learning. It’s exciting to contemplate how AI and cognitive computing will redefine KM.

Intelligent Content

Enhancing content to make it more intelligent is also top of mind for Mark Gross, President, Data Conversion Laboratory (DCL). For him, it’s important to realize how AI, ML, and NLP contribute to the new intelligence quotient of content and data. Calling it the “wizard behind the curtain,” DCL’s AI technologies result in data that becomes structured and curated even when it starts out as unstructured. The training sets that DCL has developed over many years amplify its ability to use technology to read the words in a document, understand the context, and structure the text so that it is no longer hidden in native “free-form” content.

Gross knows that information within organizations isn’t always easily accessible. It might even still be on paper, or electronically in Word or PDF. Embedded in these files are elements, such as tables, charts, diagrams, non-English

characters, or scientific formulas. AI technologies enable DCL to turn what was previously not a cost-effective project into actionable information. Gross cites the U.S. Patent and Trademark Office (USPTO), with its millions of patent applications containing multiple formats, as a shining example of how to use AI technologies to process these applications and clear backlogs.

Insights from Cognitive Search

The influence of Google on search is enormous. For many people, the words are synonymous. Thus, when employees need information from within their enterprise, they expect the search functionality to be identical to Google. For many reasons, it usually isn’t. Michael Cizmar, Founder & Managing Director, MC+A, thinks that’s going to change. He sees a paradigm shift in search technologies, one that revolves around insight engine maturity.

Cizmar identifies five levels of maturity. Level 0 is where most companies are. For them, search is not a priority. At level 1, companies realize they have an access problem, with no formal governance program and no KPIs. Level 2, Relevancy Defined, initiates governance, maps the user journey, and begins to identify KPIs. Relevancy in Context is maturity level 3. This employs ML-based relevancy, NLP query support, and an emergent search interface. The highest maturity level, level 4, occurs when there’s an auto-classification of content, predictive search, and True Knowledge Assist. To help you on your insight engine journey, Cizmar suggests a 7-step approach.

Sean Coleman, CTO and Chief Customer Officer, BA Insight, also has concerns about the value of search. Citing Stephen Covey’s landmark book, *The 7 Habits of Highly Effective People* (Simon & Schuster, anniversary edition, 2013), Coleman concentrates on the second habit, “Begin with the end in mind.” To do this, Coleman shares what he believes are the five goals you should strive to achieve when implementing search.

His first goal is to have searches return fast results. Second, users should find relevant information on the first page, no scrolling down to locate it. Third, users shouldn’t have to click more than three times. Fourth, users shouldn’t actually have to perform a search, since information should proactively flow to them. Finally, 95% of users should succeed in their searches. He also doesn’t particularly care whether search is known as cognitive

search, enterprise search, insight engines, or something else—his five goals are still key.

Living in a Digital World

AI technologies have a place in our digital world that’s not just about search or written documents. Call centers face challenges unique to the spoken word. We’ve definitely moved beyond early phone systems and should make the most of newer technologies. Here again, bringing AI technologies to bear can optimize performance. Chris Caile, Principal Product Manager for Nuance, identifies four key principles that he considers “central to a truly great caller experience.” Digital, he notes, has not killed IVR (Interactive Voice Response) systems.

To deliver value, the IVR should enjoy an intuitive, personalized first point of contact. At the very least, Automatic Number Identification (ANI) tells you the identity of your caller. Voice biometrics is another possibility. Once connected, your IVR should deliver natural, conversational interactions. Callers’ expectations are high; they’ve become accustomed to talking to machines such as home assistants and smart speakers using natural language. The third principle is to anticipate a caller’s need. AI excels at this. Finally, connect your IVR and digital channels. Omni-channel customer engagement is becoming the norm.

Cognitive and AI technologies are an integral part of customer engagement, according to Anand Subramaniam, Senior Vice President, WW Marketing, eGain. Using AI to improve customer service involves a system that understands and answers the customer’s concerns. Virtual assistants should use reasoning technologies to guide customers through complex interactions. Preferring AI reasoning to rule-based systems aids in decision making. Technology isn’t everything, however. Complying with the regulatory environment is also crucial to success.

Do It Right

To do AI right—and possibly to appease the CEO demanding AI right now—consider the myriad of responses to implementing AI technologies in a meaningful, sustainable fashion. Be clear on the ultimate goal and the immediate problem crying out for a solution. Consider which of the AI technologies, or which combination of them, provide the best answer to your individual situation. Then you’ll know you’ve done AI right. ■

AI Guidelines for Businesses: Using AI in Your Own Company

By Daniel Fallmann, CEO, Mindbreeze

Artificial intelligence (AI) is one—if not the—key technology of our decade. Technological advances in this field are not only fundamentally changing our economies, industries and markets, but are also exerting enormous influence on traditional business practices, many of which will disappear, while others will be transformed or completely reinvented.

AI Beyond Autonomous Driving and Language Assistants

Behind the notion of “AI” lies a multitude of different methods, procedures, and technologies, which have a variety of definitions.

Gartner, for instance, defines the term this way: “Artificial intelligence (AI) applies advanced analysis and logic-based techniques, including machine learning, to interpret events, support and automate decisions, and to take actions.” By comparison, Forrester defines the term as follows: “Artificial intelligence (AI) has the potential to fundamentally remake the nature of firms, employment, and how work gets done. Get insights into the development of AI.”

But what exactly does the use of AI mean for a company’s management, for an employee in a call center, for a doctor, or for scientific research staff?

The areas of application are so diverse that it is often difficult to find the right point of departure. After all, AI technologies can be used in virtually all sectors of a company to make processes more efficient or to find new approaches for mastering tasks.

In order for AI to be successfully integrated and to achieve the desired success, companies should carefully consider the initial steps. With more than 14 years of experience in implementing information insight and cognitive services based on innovative technologies, we have discovered that starting with clearly defined use cases in the relevant specialist department is the key step that yields the best results and has a long-term positive impact on the entire project.

Consider these simple steps:

1. Identify and Determine the Use Case

Companies are looking into the use of AI because it is a core element for boosting efficiency and is an integral part of digitalization that, in short, can simply no longer be ignored. Too often, we encounter a situation where a company commissions an evaluation without thinking specifically about which problems will be solved or which business processes need to be optimized. Instead of being blinded by the vast capabilities

“The use of AI opens up wide-ranging possibilities for companies, which can positively impact the strategic and operative position while at the same time generating competitive advantages.”

of the solutions offered, consideration should be given to where (which use cases) AI could be applied and which common pain points can be eliminated by using it.

The first step towards sustainable implementation should therefore be to pinpoint which problems need to be solved within the company, which can be improved through the use of AI, and which objectives need to be achieved. It will quickly become apparent that there are several possible points of departure for implementation. Dealing with all of them at the same time, however, is not particularly efficient. It makes sense to select a specific use case from a particular department to start with and then to consider how it can be optimized through the use of AI.

2. Determine and Define the Success Criteria

Once the correct use case has been identified, it is also easier to determine the next



Daniel Fallmann

Daniel Fallmann founded Mindbreeze in 2005 at the age of 23, after he finished his studies in computer science. As Mindbreeze’s CEO he is a living example of high quality and innovation standards. His passion for enterprise search and machine learning in a big data environment has fascinated not only the Mindbreeze employees but also their customers.

steps, such as defining the success criteria. Not all available solutions are suitable for every use case. For this reason, a few parameters, so-called “success criteria,” need to be defined in advance.

These criteria are:

- ◆ Business needs
- ◆ Data sources and data quality
- ◆ Semantic relationships and extraction
- ◆ ROI calculation

Business needs refer to the discrepancy between the actual situation (real-world practice) and the target situation when the solution has been optimally implemented. After the specific discrepancies have been formulated and the implementation requirements have been defined, companies

need to know what data they need from which data sources in order to achieve the desired objectives. After all, businesses possess a tremendous wealth of data that is stored in countless different sources—or perhaps has even been forgotten. Sensor/machine data, big data, documents, Internet, SharePoint, and historical company data are just a few examples. In order to extract information from this existing data and to make it usable in the form of knowledge, the relationships between the various pieces of information must be extracted and models need to be constructed so that further information can then be correctly interpreted and interlinked.

In terms of impact monitoring, it is necessary to define meaningful KPIs (Key Performance Indicators) in advance that can be understood by both employees and other stakeholders. These provide the hard facts used to measure success and serve as the basis for the ROI calculation.

“One factor that should not be underestimated when it comes to successfully implementing AI-based solutions is the quality of the existing company data. Many challenges can be solved with the help of AI, as long as the required data is of reasonable quality.”

3. Hands-on Testing With Company Data

A proof of concept (PoC) is an important milestone for implementing AI solutions. It provides the foundation for further decisions and should help separate the wheat from the chaff among the vendors. In our experience, it doesn't make much sense to try to establish the PoC simply by presenting a set of slides. The customer should use their own data to test whether the identified requirements can actually be met using the solution. At the same time, a PoC using a company's own data makes it possible to identify problems at an early stage, and the results can then be integrated seamlessly into live operation.

One factor that should not be underestimated when it comes to successfully implementing AI-based solutions is the quality of the existing company data. Many challenges can be solved with the help of AI, as long as the required data is of reasonable quality. “Data garbage” can significantly complicate the use of machine learning and AI. For example, if the data is incomplete, present in multiple versions in different data sources, contains inconsistencies, spelling, punctuation or other general errors, it will be difficult to obtain correct and relevant results (garbage-in-garbage-out principle).

To avoid this, the existing data in the various data sources needs to undergo a more detailed examination beforehand:

View and understand the data: It is necessary at this point to consider which data sources are available in the company and whether all these data sources are relevant for the planned PoC. It may prove useful to first work with and learn from the most relevant data sources. When selecting the solution, care should be taken to ensure that the pool of data sources can be expanded as required.

Data cleansing: The existing data has to be subjected to a critical examination to determine whether all existing file shares are still needed.

Process and link the data appropriately: The better the chosen product is, the less manual effort this step will entail. The ability to automatically recognize essential correlations and, most importantly, to render them visible to the user together with the possibility for providing feedback is a key factor to consider when selecting the right tool.

The steps described above can be carried out in a very short time but are immensely important for the quality of the results.

4. A Critical Step: Involve the Users

If the AI solution in the PoC for the selected use case looks viable and convincing from a technical point of view, the experts from the relevant department need to become involved. Therefore, in keeping with the principle of continuous change management, employees should become involved as early as this phase. They know their processes best and are in the best position to judge where the solution still needs to be improved, and they can also provide valuable input—especially when it comes to training the AI solution. Only through continuous use and active feedback can AI learn new things, be trained, deliver more accurate results, and consequently support the day-to-day work. Quite often, very minor optimization measures—such as fine-tuning the relevance model—yield enormous added value. Feedback from key users is essential for a successful subsequent roll-out as well as for acceptance, because they act as multipliers in the company.

In this context, the different deployment options can be considered and adapted to the relevant business needs. AI solutions are usually available in different variations. To find the right option, companies should find out WHERE—in which data sources—they store the majority of their information.

◆ **Appliance:** The appliance version involves high-performance hardware with perfectly tailored software. This is integrated directly, without any connection to external systems (Internet, manufacturer, provider) into the company's internal data center and indexes the data from the connected data sources. This option is favored by companies that do not have any data stored outside the company or that want to use the AI solution to process data that is particularly sensitive.

◆ **SaaS:** For companies that manage a great deal of data in cloud solutions (SharePoint Online, Office 365, Salesforce), a SaaS solution is the ideal option. The AI solution in this case is located in an external data center.

◆ **Hybrid:** With a hybrid solution, the data stored in the cloud is merged with the data from the company's own data sources.

The AI solution can be operated either in the cloud or on-prem.

5. Validate the ROI Calculation

If the test results are satisfactory, the transition to real operation can begin. All settings can be transferred directly from the PoC, and the ROI calculation that was created at the start can be validated. Once the project has been successfully launched, the AI system can be rolled out across other departments or business areas until it is ultimately deployed throughout the entire company, thus transforming all processes (business process transformation).

When used correctly, AI can dramatically improve a company's performance by automating recurring business processes, increasing understanding of the customer, improving efficiency, and boosting employee satisfaction.

A checklist/cheat sheet for implementing AI-based processes:

- ◆ What is the specific problem—what is the use case?
- ◆ What is the objective—what do we want to achieve with it?
- ◆ Where is the data located—in which data sources?
- ◆ What do the users say—what needs to be improved?
- ◆ Have we achieved our targets—what is the return on investment?

Nowadays there are opportunities for optimization along the entire value chain. In some specific application areas, AI is already more powerful and efficient than humans. This is especially true for routine tasks such as distributing or classifying documents, analyzing large amounts of data, or recognizing patterns. Routine tasks include things like visual recognition, speech recognition, natural language processing, and translation. But even tasks that require a certain degree of logical understanding and learning, as well as tasks that are carried out by experts, can be supported in the future through the use of AI. The final authority for decisions, however, is still the human being.

My personal conclusion: AI in companies, when well-thought-through and intelligently implemented, can be viewed as a real game changer and a competitive advantage, both from an economic perspective as well as from a personal perspective. ■

Understand. Anticipate. Improve.

How Cognitive Computing Is Revolutionizing Knowledge Management

By Kelly Koelliker, Director of Content Marketing, Verint

For decades, organizations have tried to unlock the collective knowledge contained within their people and systems. And the challenge is getting harder, since every year, massive amounts of additional information are created for people to share. We've reached a point at which individuals are unable to consume, understand, or even find half the information that is available to them.

Artificial intelligence (AI) and cognitive computing technologies have emerged as a disruptive and innovative force across many industries. These new technologies are now helping to simplify, modernize, and automate knowledge solutions. By bringing together traditional knowledge management tools with advanced computing intelligence, organizations can vastly simplify the process of creating, finding, and improving knowledge. This powerful combination, which we call *cognitive knowledge*, is a more natural and effective way for businesses to connect people to knowledge.

Advancements in AI now enable people and machines to interact more naturally to extend human expertise and cognition. In fact, AI and cognitive computing can augment your knowledge management program by helping you understand what people are searching for, anticipate follow-on requests, and continually improve the information that you provide in the future. Let's take a closer look at each activity.

Understand

The most fundamental component to a successful knowledge management application

is the ability to easily search for information. AI can understand the nuances in how people ask questions and search for answers. Like our brains, AI disambiguates words and concepts based on context. Sophisticated AI algorithms leverage mathematics underpinned by an extensive body of linguistic data. Similar to the human experience, advanced knowledge management systems can understand the correct meaning of search terms, even if they are phrased in different ways. These applications can analyze content semantically to determine contextual relevance to a user's request, pinpointing solutions more accurately than non-cognitive approaches.

Anticipate

Beyond simply understanding a user's current request, cognitive computing can anticipate the answers users may need and predict what they might ask next. AI-based knowledge management is both proactive and predictive, making inferences based on the conceptual understanding of the knowledge base to deliver useful information when and where people need it.

Similar to human memories, these applications cluster related concepts and recall this knowledge to predict next questions. Essentially, knowledge applications can mimic a person's train of thought, answering questions before they are even asked. For example, an initial inquiry about the symptoms of a medical condition might trigger suggestions on relevant treatment options, medications, and lifestyle changes.



Kelly Koelliker

Kelly Koelliker is director of content marketing at Verint with a focus on contact center workforce engagement solutions. With more than 15 years of marketing and sales experience, her expertise in the customer service industry covers such fast-evolving categories as knowledge management, natural language search and CRM.

Improve

The third piece of the puzzle is implementing a cycle of constant improvement. By enhancing content automatically through continued use, less manual administration is required. AI-based systems have the ability to incorporate both implicit and explicit knowledge improvements to eliminate administrative heavy lifting. Instead of building taxonomies or requiring specialized learning, these applications can automatically build models of conceptual understanding, and continuously refine these models through prolonged use and administration.

These applications can provide insight into search accuracy, usage patterns, and knowledge gaps to help administrators improve content through the entire knowledge lifecycle and better ensure information is always fit for purpose.

Cognitive computing can also improve knowledge quality by monitoring real-time data streams from digital sources, such as email correspondence and social media feeds, and uncover trends as they arise. For example, it might reveal a sudden upsurge in questions about how to use a new product feature. This can help improve the responsiveness of knowledge workers, enabling them to avert potential issues before they disrupt your business.

To sum it up, cognitive knowledge mimics human thought processes and models expertise to answer questions accurately and consistently. This cognitive approach enables a new generation of employees to provide fast, accurate, and effective advice to help ensure customer success, fuel agent performance, and maximize employee productivity. Behind every exceptional digital experience is a truly great technology enabling incredible things to come together. AI and cloud computing are redefining the world of knowledge management for an exceptional customer experience with less effort. ■

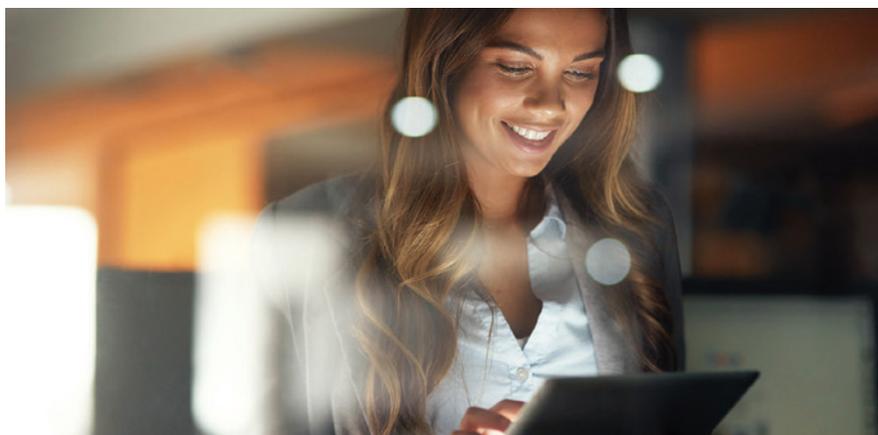
Contact Us:

1-800-4VERINT

www.verint.com/engagement

www.verintblog.com

© 2019 Verint Systems Inc. All rights reserved worldwide.



3 Things to Know Before Starting Your AI Journey

By Scott Parker, Director of Product Marketing, Sinequa

AI-Powered Search Engines—referred to as “Insight Engines” by Gartner and “Cognitive Search” by Forrester—can deliver significant value to organizations these days, provided certain risks are avoided.

As the name suggests, these engines automatically analyze enterprise content and data to surface relevant information and insights. They automatically surface relevant information for users within the context of their work and, depending on the specific use case, result in better decision making, superior customer service, more effective management, and significantly improved performance overall.

In recent years, technology to simulate human intelligence processes, or artificial intelligence (AI), has become an increasingly useful component in these systems for automating and in some cases scaling certain types of analysis. The science of getting a computer to perform specific tasks without using explicit instructions, better known as machine learning (ML), has produced so-called “supervised” learning algorithms that take labeled data sets (training sets) to automatically detect patterns and generate models that can be used to label new unlabeled data sets. Deep learning, a subfield of machine learning, involves the training and application of advanced mathematical structures referred to as artificial neural networks. Deep learning enables advanced processing of free text content and opens the door to newer and broader types of artificial intelligence such as open question answering and user intent recognition.

That all seems pretty compelling but maybe not so straightforward. In fact, according to a recent survey conducted by IDC, 25% of organizations worldwide that are already using AI solutions report up to a 50% failure rate. A lack of skilled staff and unrealistic expectations were identified as the top reasons for failure. [IDC global survey of 2,473 organizations, May 2019.]

As an independent software vendor providing an AI-Powered Insight Engine, Sinequa has direct experience with many AI initiatives. Below are the best practices we recommend following to ensure success.

Identify and Focus on the Right Use Case

First and foremost, “doing AI” is not the end game, addressing a valuable use case is. If a use case requires AI, validating that it is compelling to the business is key (ROI, usability, etc.). At Sinequa we have created a dedicated AI department to work with cus-

tomers on their AI-related projects. Experts within this department analyze, jointly with customers, the highest-value use cases for AI usage, the data they have (while checking its adequacy), define metrics that are used to measure results, and use the results to prioritize and design customer road maps.

It is critical that the use cases be grounded in reality, meaning they should be well within the capabilities of today’s technology, and the use cases should merit an AI-powered approach. Many use cases can be accomplished in a simple and straightforward manner without the need for artificial intelligence technology. However, there are certain use cases that can only be realistically and cost-effectively accomplished with AI technology, like auto-classifying large volumes of confidential documents or auto-routing large volumes of support tickets based on the case details, for example.

Implement the Right Methodology

Assuming a solid business case has been established, an effective methodology for AI projects begins with *Data Acquisition*

and Preparation. Data acquisition entails establishing data availability, connectivity, and security, whereas data preparation involves steps such as cleaning, normalizing, labeling, and enriching. Administrative tool support for this step is something to look for when considering an Insight Engine for your AI project. The next step is *Model Building*, which entails model selection, hyperparameter tuning, and feature selection. Again, adequate tool support is essential to make this as simple and straightforward to execute as possible. The third step is *Deployment & Validation*, which involves frequent model re-training and application, systematic validation of output quality, and integration of the model’s prediction within third-party systems. The final step is *Active Learning & Tuning*, which involves providing labeling tools to subject matter experts, managing training sets and test sets, mitigating bias



Scott Parker

Scott Parker has a deep history in the enterprise software business helping organizations to succeed by becoming information-driven. He was a founding member of the IBM Watson Group where his team engaged clients to explore the art

of the possible around augmenting human intelligence in all areas of the enterprise. Parker joined Sinequa at the beginning of 2017 and he works relentlessly to make sure stakeholders across knowledge-intensive industries understand the real and potential benefits of becoming information-driven.

and edge cases, as well as setting up model versioning and lifecycle support.

Use a Proven and Complete Technology Platform

Whatever technology is chosen to support your AI projects should have been forged by experience with other organizations with similar goals and complex projects. Specifically, the technology should provide hardened capabilities such as enterprise-grade security, extensive linguistic and natural language processing capabilities, integrated machine learning, and well-designed user experience.

In order to realize value quickly without sacrificing context or quality, the various technology components must be pre-integrated and cohesively designed to form a unified,

“In order to realize value quickly without sacrificing context or quality, the various technology components must be pre-integrated and cohesively designed to form a unified, end-to-end solution. By adopting a unified solution, organizations shield themselves from the cost and complexity of IT integration.”

end-to-end solution. By adopting a unified solution, organizations shield themselves from the cost and complexity of IT integration. They also benefit from a development paradigm that enables implementation and administration through configuration, scalable to even the most complex environments. Adopting a complete solution also means these organizations do not have to purchase a point solution whenever new business requirements arise. Instead, they can leverage a unified, end-to-end platform to configure applications quickly that are specifically aligned with business goals.

Conclusion

As many organizations have experienced, adhering to these best practices helps move projects that use AI technology from the experimental phase into production and greatly improves the odds of achieving the predicted return on investment. ■

From “Searching” to “Finding”: How AI is Unlocking the Power of Unstructured Data

By Kamran Khan, Managing Director, Search & Content Analytics, Accenture

Unstructured data, which comprises almost 80% of any enterprise’s data, holds untapped value when it comes to addressing challenges and embracing opportunities. Extracting valuable insights from unstructured data sources has traditionally been difficult for businesses because it involves complex and time-consuming data analytics processes. However, with the help of natural language processing (NLP) and machine learning (ML), this is changing fast.

If businesses could take full advantage of the value this information holds as and when they need it, they would be able to resolve and even prevent issues more efficiently across the whole enterprise.

However, unlike structured data such as tables and spreadsheets, which have long been put to good use within enterprises, unstructured data is much more difficult to leverage and a lot harder to proactively analyze. Until recently.

In recent years, pragmatic AI—a collection of multiple technologies that encompasses deep learning, NLP, search, machine learning, intelligent sensors, and robotics has become a critical driver of enterprise evolution. It powers a variety of intelligent tools that are transforming the data supply chain for better insight discovery. Three prominent AI capabilities are driving a lion’s share of the change: Internet of Things (IoT), computer vision, and document understanding.

Of the three, document understanding is helping to make it possible for organizations to extract valuable insights from untapped, unstructured data sources. It combines NLP and ML to help gain insights into human-generated, natural language unstructured text. This should play an increasingly influential role in the future of enterprise transformation, as the number of unstructured sources within organizations will likely increase.

Natural Language Processing Gains Momentum

The practical applications of NLP make it an ideal tool for businesses looking to leverage

insights from unstructured data to transform their operations.

As a result, NLP has been going through a transformation of its own, with two main factors driving this change.

First, as business users become more accustomed to search engines like Google and digital assistants in their personal lives, they begin to expect the same knowledge-seeking experience at work. This is fueling high-performing enterprise NLP applications that can understand and respond to natural language queries.

Second, NLP no longer relies on manually written rules alone. For greater automation, scalability and accuracy, NLP is now being paired with ML and enabling tools such as document understanding applications.

This shift also explains why the analysis of unstructured data is increasingly moving away from reactive searches and towards the proactive generation of insights to feed existing or anticipated business needs. Document understanding applications that combine NLP and ML deliver just that.

Document Understanding Drives Productivity

By combining search and analytics with pragmatic AI technologies like NLP and ML, document understanding automatically extracts relevant information from unstructured data sources, saving businesses the time and resources needed to search manually.

As these applications further develop, they can deliver advanced actionable insights to improve business processes and customer experience. In fact, businesses across many industries have started to apply document understanding to help surface insights, including:

- ◆ **Legal departments**—Reducing risks by automatically analyzing legal contracts for specific “red-flag” terms.
- ◆ **Government agencies**—Analyzing digitized incoming mail to route relevant letters to the right departments, eliminating manual effort and saving hundreds of thousands of agent hours.



Kamran Khan

Kamran Khan has been developing, supporting, selling and managing in the computer software/services industry for more than 25 years. In 2005, he founded Search Technologies, a global IT services company dedicated to implementing and supporting search and big data solutions. He served as president and CEO until their acquisition by Accenture in July 2017, and now is a managing director of the Search & Content Analytics Group within Accenture Applied Intelligence.

He is a frequent contributor to CMSWire.com and a repeat keynote speaker at KMWorld annual conferences in addition to being interviewed for many other publications such as *Forbes*.

- ◆ **Recruiting**—Taking on rote tasks like sifting through millions of resumes and automatically matching CVs to job postings.
- ◆ **Banks and financial services**—Automatically cross-analyzing loans or mortgages with the borrowers’ profiles from multiple independent sources to deliver better customer experience and engagement.
- ◆ **Storage optimization**—Using automated business rules to identify the appropriate action to take with documents stored in expensive on-premise storage—whether to move to lower-cost storage, delete, or archive. ML can also accurately and quickly detect duplicates, allowing for storage cost savings as well as a 360-degree view of enterprise data.

Poised to Reach New Potential

With the increasing range of pragmatic AI solutions available, from open source frameworks and evolving vendors to cloud-based APIs, enterprises stand to benefit more than ever from this ecosystem. They now have the flexibility to integrate appropriate approaches and technologies for their use cases.

While NLP is not perfect, it is being consistently enhanced. And the ML algorithms supporting NLP are seeing significant advances with industry giants like Google, Microsoft, and Amazon making strides to improve accuracy.

We’re also leveraging our own technology assets at Accenture to orchestrate different components of NLP applications, making them easily maintainable and scalable using both custom and ready-built algorithms. This means that NLP and ML are slowly gaining maturity, helping businesses to use document understanding to tackle increasingly complex challenges and finally begin to unlock the full potential of unstructured data.

Now that you know more about document understanding, what potential use cases do you see within your organization? What value will it help you unlock? ■

AI and the Building Blocks of Intelligent Content

By Mark Gross, President, Data Conversion Laboratory

The data, information, and analytics economy runs on well-curated, structured data. No matter your industry—having good curated data and content is critical. It's increasingly important as more data and content are generated. Intelligent tools to sift through content are more robust and at the same time, more "needy." That means modern technology platforms, systems, and even content consumers require well-structured data and content to perform well. As most artificial intelligence (AI) practitioners state—"nothing starts without good data."

This article explores the use of AI, machine learning (ML), and natural language processing (NLP) in the construction of intelligent data and content.

The terms have exploded recently in the business vernacular but the concepts trace back to the 1920s when the [Austrian engineer Gustav Tauschek obtained the first patent for his "Reading Machine."](#) Optical character recognition (OCR) is taken for granted today and the process is commoditized. However, early OCR systems needed to be trained with images of each character and worked on one font at a time. Sounds like machine learning, right? Likewise, the field of computer vision is a precursor to AI image recognition and traces its origins back to the 1960s. While there were developments in all these areas throughout the 20th century, and earlier, it's come of age in the past 15 to 20 years due to the confluence of fast computers, inexpensive large-scale storage, the Internet, and data collections. AI, specifically ML, needs large amounts of training data, and fast computers to process it all—today we have that.

How are AI, ML, and NLP Relevant to Data Structure?

How do businesses go beyond the hype of these terms to ensure they implement the right technology and truly get valid ROI?

Many organizations have extensive and valuable data and content buried in paper, PDF, and Word files. The content is not structured and not necessarily even digitized. Unraveling this, especially if the content includes complex tables, charts, figures, foreign characters, chemical formulae, etc., was impractical. Until now. Advances in technology and in AI make intelligent data and automated content

structure possible where it just wasn't feasible before.

Technology enables multidimensions with what was considered "flat content." You can't do anything with flat content—no search, no filters, no interactivity, no related content. But if you properly structure data and content, you apply a layer of semantic intelligence that benefits the downstream delivery and consumption of that information. In other words, you can deliver or receive information with precision.

At Data Conversion Laboratory, we say that AI and related technology enables organizations to revisit high-value, but pre-

"AI empowers products and ideas that were previously impossible."

viously impractically expensive, projects. Common attributes of data and content that could benefit from AI and related technologies include:

- ◆ **Digitized without structure**—Scanning alone, with "dirty" OCR, is a good first step to preserve documents, but the results are image-based PDFs, which are not easily searchable. Modern AI, specifically NLP, can extract intelligence from that previously digitized content. Organizations can achieve more than what was feasible before!
- ◆ **Complex content or data**—In the past, variable data (or content) types with special characters, math, chemical formulae, etc., were "digitized" as images. That means filtered search and data analysis could not be performed on or with this type of information. New levels of accuracy are now possible with computer vision.
- ◆ **Security and automation**—New AI techniques provide capabilities that were previously impossible in a manual or semi-automated process. Now, cost-effective solutions exist to deal with confidential or sensitive data.

The Wizard Behind the Curtain

Some people don't really care how data or content is structured. But it's important to understand *if and how* you are using intelligence and technology to structure data or content (or if your vendor is!). DCL uses onshore staff with top technology to create structure where it didn't exist before. The technology has evolved over the years, thanks to our work with structured markup languages. We have built large training sets that enable us to use AI / NLP and generate accurate and rich content.

Let's look at a simple example. The following text is an example of a simple paragraph referencing mechanical parts:

FIG. 6 illustrates a diagram of the signal adaptive pre-filter 1200 and motion detector 1300 section within the segmented temporal processor 1400.

NLP looks at words, the order of words, and neighboring words and is able to discern what exactly is a "part":

```
<para> FIG. 6 illustrates a diagram of the
<part-name>signal adaptive pre-filter</part-name>
<part-number>1200</part-number>
and <part-name>motion detector</part-name>
<part-number>1300</part-number>
```

```
section within the <part-name>segmented
temporal processor</part-name> <part-number>1400</part-number></para>
```

This is possible because DCL has developed extensive AI training sets over the decades. The combination of NLP, computer vision, and automation enable the computer to "read," "understand," and contextually structure complex technical text buried in free-form content.

Real-World Application: The United States Patent and Trademark Office (USPTO)

The USPTO processes millions of trademark and patent applications. The information is dense and a combination of unstructured text, images, math, metadata, and more. Patent examiners required a system that allows them to search information in patent applications. They required a process to take unstructured, confidential information and structure it—with **zero human intervention**.

Explore the details and complexities of this project by downloading the white paper, ["Lights-Out Automation: Using AI to Create Structured Data From Static Documents."](#) ■

A Best Practice Approach to Insight Engines: 5 Levels of Insight Engine Maturity

By Michael Cizmar, Founder & Managing Director, MC+A

Enterprise search projects start with intentions to provide ‘Google for our organization’ but too often fail to deliver on that promise. In our experience, these projects fail due to a lack of sustained effort and governance. The commercialization of next-generation search technologies allows you to fulfill this promise if you take a systematic approach to implementation.

Embracing the Era of Insights for Search

We are witnessing a paradigm shift in technology where it is possible to deliver unique search experiences to every user at every touchpoint. Driving this shift is a collection of technologies that industry analysts refer to as Insight Engines. Insight Engines, as defined by Gartner, serve to “augment search technology with artificial intelligence to deliver insights—in context and using various modalities—derived

from the full range of enterprise content and data.” They embed AI technologies within traditional search infrastructures. This embedded AI makes them aware of the context from which a user is searching while understanding where users (should) want to go.

Insight Engines are a collection of component technologies rolled into platforms, rather than a singular technology. Subcomponent technologies to an Insight Engine can include Natural Language Processing, Machine Learning, and Artificial Intelligence. These technologies are game-changing, but their impact is realized when implemented through a systematic approach that is focused on how they can improve a user’s search experience. When focused on a user, their journey, and the fact that searching is not what they intended to do when visiting your website or using your application, you deliver better outcomes. Better outcomes for users drive improvements to all other business metrics.



Michael Cizmar

Michael Cizmar is the managing director and founder of MC+A with over 20 years of experience in software design, consulting, and technology. In addition to leading the MC+A sales team, Cizmar also oversees the

company’s Insight Engine practice and is deeply involved in the evangelization of next-generation search technologies, including NLP, ML, NER, and AI.

A Maturity Model for Delivering Insights

Improving your search maturity is a vital, multi-step process. It starts by understanding where your organization is at on the maturity ladder below.

Understanding your current maturity level and where you aspire to go, allows you to road map a path forward. Briefly, these are the steps:

- Step 1**—Form an internal committee to lead the initiative
- Step 2**—Build a unified index of all the content that your users are interested in
- Step 3**—Map the search experience to the user journey
- Step 4**—Develop KPIs
- Step 5**—Implement ML-based relevancy adjusting relevancy based on the use of the system
- Step 6**—Implement user intent detection and provides answers
- Step 7**—Implement intelligent recommendations

When working through these steps, you should also review trends in consumer platforms (Google, Facebook, Amazon) to understand what user expectations could be when interacting with your website or application.

The steps described above provide iterative benefits and should produce feedback to guide next steps. Insight Engines allow you to focus on the user and deliver incredible experiences by understanding intent and delivering intelligent recommendations. For more details on our maturity model and its stepwise approach, go to www.mcplusa.com/maturity-model/.

MC+A has implementation expertise deploying Insight Engines across many verticals and use cases and is interested in discussing how we can help you deploy them for your users. ■

Level	Short Name	Described By
0	Most Companies	<ul style="list-style-type: none"> • App-based or departmental • Searching for documents user knows exist • Unaware of industry capabilities • Search not a strategic priority
1	Realization of Access Problems	<ul style="list-style-type: none"> • A unified index • No formal governance • No Key Performance Indicators (KPIs)
2	Relevancy Defined	<ul style="list-style-type: none"> • Governance initiated • Consideration/mapping user journey • Initial KPIs identified • Active search tuning to KPIs
3	Relevancy in Context	<ul style="list-style-type: none"> • ML-based relevancy (Learn to Rank) • NLP/NER query support • Emergent search interface
4	Insight Engine Achieved	<ul style="list-style-type: none"> • Auto classification of content • Predictive search • True knowledge assist

IVRs and AI, Unite!

By **Chris Caile**, Senior Principal Product Marketing Manager, Nuance

While interactive voice response systems (IVRs) have been invaluable in reducing contact center costs, we need to be honest: not many are delivering experiences that meet consumer expectations. It's no surprise given the rise of digital channels.

But digital hasn't killed the IVR—it's pushing it to be better. Customers expect the same seamless, intuitive experiences they get in the digital world everywhere they engage your organization—especially in the IVR. The phone is more important than ever for your customers: it's an escalation point if digital channels don't deliver what they need, it's the go-to channel for anyone who is 'on the go' and can't easily navigate a website or an app, and it's the top spot for complex or sensitive issues that require human assistance.

So, at your customers' critical moments of need, are they greeted by an outdated touch-tone phone system from 1998 with awkward menu options? Or an intuitive, conversational and effortless experience that quickly addresses what they need?

Four AI Solutions for a Modern IVR

We've identified four key principles that are central to a truly great caller experience. Each one involves the power of artificial intelligence (AI). AI is key to reinventing IVR because it moves it beyond simple voice recognition to a deeper level of understanding and interaction.

1. Create an Intuitive, Personalized First Point of Contact

Imagine if you and your closest friends had to re-introduce yourselves every time you met. Or every time you ordered a coffee from your local barista, you had to verify your identity. It sounds ridiculous—but that's basically what outdated IVRs are putting even your most loyal customers through.

There are far better ways to offer personalized first contact experiences. At the simplest level Automatic Number Identification (ANI) eliminates caller identification prompts completely. Instead, your IVR automatically identifies inbound phone numbers and matches them against your customer data to work out who's calling before the interaction even begins. Sounds easy, but many organizations haven't taken the time to do this most basic step.

There's also a far better way to handle authentication painlessly: voice biometrics. With voice biometrics, your voice is your password. The IVR automatically verifies your identity by comparing your voice against a "voiceprint" you created during a simple one-time setup process. It's not just a more pleasant experience—it's far more efficient and secure. On average, organizations have

seen call durations shortened by an average of 42 seconds, freeing agents to spend time on more complex issues.

Combining both ANI matching and biometric authentication will make your IVR a joy to call.

2. Deliver Natural, Conversational Interactions

Talking to an IVR isn't new. For years, customers have been able to speak a word or two to make a menu selection or get to an agent. But the IVR typically relied on the caller repeating a pre-programmed set of words—not

"Putting AI to work in the IVR enables organizations to predict the reason for a call and automatically deliver a highly personalized experience tailored just for them, based on their unique situation and past history."

asking questions or exchanging dialog.

Those kinds of one-way interactions don't cut it anymore, especially with today's home assistants and smart speakers. Today's callers expect your IVR to listen and respond naturally as if they're talking to a human. And, thanks to huge breakthroughs in AI and NLU, today's modern IVRs can have a real, two-way conversation. A true dialog with your customers without having to master a crazy menu prompt. Imagine your callers being greeted with "Hello Bob, thanks for calling. How can I help you today?" and then the IVR helping resolve the issue.

3. Anticipate Your Caller's Needs

We don't call a company because we are looking to talk with someone. We call with questions. And we want to spend as little time as possible on the phone. Forward-thinking companies are using the latest advances in AI and predictive capabilities to better anticipate why customers call and get them to the right resolution faster.

Putting AI to work in the IVR enables organizations to predict the reason for a call and automatically deliver a highly personalized experience tailored just for them, based on their unique situation and past history.

Suppose a cable subscriber spends time researching programming packages before calling. The IVR can recognize that activity and greet the caller with a personalized message: "I



Chris Caile

Chris Caile joined Nuance in September 2015 as senior solutions marketing manager for Nuance Conversational IVR (Interactive Voice Response). Before joining Nuance, Caile worked in various marketing and sales support positions at Microsoft and Motorola and has over 20 years' experience in the high tech industry. Caile holds a bachelor's degree in business administration from Illinois State University with minors in mathematics and economics.

see you were browsing cable packages online earlier today. Are you calling to add a station or change your programming lineup?"

By analyzing information such as browsing history and past transaction data, you can anticipate customer needs before they even say "hello." And the best IVRs don't just resolve the customer's issue more effortlessly; they recommend the right next step for every customer, too.

4. Connect Your IVR and Digital Channels

Companies with the strongest omni-channel customer engagement strategies retain an average of 89 percent of their customers, compared to only 33 percent for companies with weak strategies. Why? Because 60 percent of customers visit your website before calling your IVR, and 36 percent continue to browse your website while talking to a representative. Wherever they start their journey, they expect to authenticate once, then move seamlessly between channels to get the answers they need.

An IVR-to-digital solution can reduce your contact center costs and improve customer satisfaction by allowing customers to transfer seamlessly from your IVR to a virtual assistant or a live chat agent.

So how does it work? Imagine a customer calls with a question about a new product they're considering. Here, the IVR offers them a choice:

"No problem, I can have a live agent help you with that—the wait time is just five minutes. Or you could chat with an agent right now via text message. Would you like to chat instead?"

Many times, customers will choose live chat to accelerate the process. Seamless connections between your IVR and digital channels creates the experiences customers expect and boosts satisfaction.

Helping Your IVR Flex

Your company's IVR delivers value, and now with [these four imperatives for improvement you can boost your IVR game and become a favorite alongside their digital brethren](#). Because everyone, even the IVR, likes to flex every now and then. ■

Key Considerations in Maximizing the Value of Cognitive Search

By Sean Coleman, CTO and Chief Customer Officer, BA Insight

I am a firm believer in *The 7 Habits of Highly Effective People*, by Stephen Covey. If you've not read this book, it is worth the time. I mention this because my focus at BA Insight is around Covey's second habit, which is, "Begin with the end in mind." Seems simple, right? Well it is, but it's also quite rare. When approaching any enterprise search project, at any phase, I always try to come back to this idea. What is success? When are we done? What does finished look like? These are all different ways of saying, "Make sure you have goals!"

I'll share with you what I see as the core goals that should be the outcome of any search system. Regardless of whether you call it Enterprise Search, Insight Engines, Site Search, or Portal Search, these goals still apply.

Goal 1: Searches Should Return Fast Results

◆ *Specifically:* Search results should be returned in an average of three seconds or less for more than 70% of searches, and they should never be longer than six seconds.

Goal 2: Users Should Find Relevant Information on the First Page

◆ *Specifically:* Users should see relevant information on the first page of search results at least 70% of the time, determined by whether they interact with an item in the search results.

Goal 3: Users Should Not Have to Click More than Three Times

◆ *Specifically:* It should take no more than three clicks for users to get to the information they are seeking for at least 90% of searches.

Goal 4: Users Should Not Have to Actually Perform a Search

◆ *Specifically:* The search system should understand users and their relationships to data and proactively suggest content to them.

Goal 5: 95% of Users Should Succeed in their Searches

◆ *Specifically:* Less than 5% of users should abandon searches. By abandon, I am referring to users who run a search and end their sessions without clicking on any filters or results.

Make no mistake, these are hard numbers to hit. At BA Insight, we focus on making each of these goals achievable. Here's how:

Goal 1: Searches Should Return Fast Results

BA Insight's SmartHub software provides the user interface components and query engine that allow for a high-performance search implementation. Its out of the box responsive/mobile-ready design means that regardless of the platform, the user experience is optimized. Its query engine supports integration with multiple search engines including Elasticsearch, Elastic Cloud, Azure Search, SharePoint Online/On premise, and SOLR (coming soon). This capability

"Tracking and reporting on metrics associated with any search system is key to understanding the success of the system."

provides organizations the flexibility to deploy the best search engine for their specific use case or infrastructure.

Goal 2: Users Should Find Relevant Information on the First Page

Relevancy is a direct result of the availability of data, quality of metadata, and intelligence wrapped around the user's query.

Our ConnectivityHub provides over 70 out of the box connectors to enterprise systems. Our experience is that most users need to search three or more sources to ensure relevant content is returned.

Our AutoClassifier provides intelligent tagging, metadata generation, and text analytics to make content findable. Content must be tagged with metadata, and metadata creation is not something you can burden end users with.

As previously mentioned, a major component of our SmartHub is its query engine. The most important key to delivering quality relevancy is to personalize results to users. SmartHub uses machine learning to adjust the relevancy of results per user, ensuring relevant information is delivered based on roles, departments, or other attributes.

In environments where finding the correct information would require users to enter complex searches, they should be able to ask questions of the search system. SmartHub fully supports Natural Language Query and can interpret questions in natural language and return relevant results based on an understanding of what users need.

Goal 3: Users Should Not Have to Click More than Three Times

This goal builds upon the previous one and points to the key aspects of a search UI. Facets, which also require metadata, must be configurable to present filtering options to users. As an example, think about the search experience in Amazon. Facets are also a key feature of our SmartHub product.

Our Smart Preview tool provides single click access to a mobile-ready instant preview of an entire document regardless of location, further reducing the clicks required to validate the results presented to users.

Goal 4: Users Should Not Have to Actually Perform a Search

This goal introduces the concept of zero search, or the ability of the system to proactively provide access to content that users consider helpful.

Our SmartHub tool allows end users to specify areas of interest and preferences and automatically delivers content to them based on their input or existing profile. For example, if a presentation that was used to train users in a particular job role has changed, then the updated version of that PPT could be proactively pushed to all users in that role.

Goal 5: 95% of Users Should Succeed in their Searches

For users to succeed in a search, they need to find what they are looking for. That much is clear. The approaches reviewed in the above goals all speak to how this overarching goal can be met. But how can this goal be measured? Tracking and reporting on metrics associated with any search system is key to understanding the success of the system. You must be able to report on at least the major goals outlined above, and hopefully on a large range of other attributes to help manage the search system over time.

Our SmartHub application comes with a built-in, fully featured analytics tracking and reporting application, ensuring organizations have access to reports and metrics across the entire search experience. These metrics are then leveraged to increase the chances of users succeeding in their searches. Prior successful search activity, combined with machine learning, automatically recommends potential searches and information to users. As users type, suggestions are delivered based on successful searches run by others. On search results pages, additional information is recommended based on what other users have found to be useful. ■

AI-Powered Customer Service: Use-Cases and Real-World Examples

By Anand Subramaniam, Senior Vice President, WW Marketing, eGain

Cognitive/AI technologies for customer engagement are white hot. No wonder professionals, who had removed AI from their resumes, are scrambling to add it back in!

As a pioneer in cognitive/AI solutions for customer service and engagement, eGain has not only developed cutting-edge technology but also proven use-cases and best practices over the last two decades. In this article, we discuss four use-cases that have already enabled blue-chip companies to transform—not just improve—customer service and engagement.

1. Understand and answer

Virtual assistants (VAs) help businesses wow customers with natural language understanding and distinctive self-service, while helping them cut costs and build brand equity. The best VAs are also multilingual and communicate in multiple modes—text-to-text, text-to-speech, speech-to-text, and speech-to-speech. Importantly, they know what they don't know. When unable to answer the customer's question, they escalate to human-assisted customer service with full context from the self-service interaction. For the VA to be able to do this out of the box without integration work, it needs to be an integral part of an omnichannel customer engagement hub which consolidates omnichannel interactions, knowledge, AI, analytics, and administration into one platform.

eGain AI clients understand and answer:

A large government organization in the UK uses the natural language capabilities of eGain Virtual Assistant™ to understand and answer questions from taxpayers, with intelligent and seamless escalation to live chat when necessary. Deployed in time to support the critical tax returns period, eGain Virtual Assistant and eGain Chat™ helped deflect 77% of their phone calls within the first six weeks of the deployment!

2. Guide search and processes

While VAs are good at answering questions of low-to-medium complexity, AI reasoning technologies can guide customers and contact center advisors through interactions of higher complexity. Reasoning can guide users to the next best steps in their search for the correct answer or in-service processes such as troubleshooting and advice. This conversational, dialog-driven guidance is based on intelligent understanding of the problems faced by the customers as well as customer service expertise drawn from the best agents.

AI reasoning applies learnings from past cases to find solutions for new ones. Learning should be derived through a curated model rather than a fully automatic one, especially when the stakes are high, to ensure the best outcomes and process adherence. Watch out for pretenders like rigid scripting and rule-based systems—they tend to put agents and customers in conversation cul-de-sacs and dead ends, especially when the customer goes off-script (which is quite common). Moreover, such legacy systems are difficult and expensive to maintain.

eGain AI clients guide:

◆ *A global bank uses eGain's patented AI reasoning technology to guide a largely novice agent pool through best-practice interactions across 11 countries, while reducing training requirements by half.*

◆ *A leading telco guides 10,000 contact center advisors and associates in 550 retail stores with eGain AI for a 37% improvement in First-Contact Resolution (FCR), 50% improvement in advisor speed to competency, and a 20% boost to their NPS (Net Promoter Score).*

◆ *A multi-play CSP reduces unwarranted "No Fault Found" handset returns and exchanges by 38% with AI-guided problem resolution in the contact center.*

◆ *A white goods giant saved \$50M per year by reducing unnecessary truck rolls through AI-guided problem resolution in the contact center.*

3. Help decide

According to an executive survey by Grove Critical Thinking, when presented with the proposition "bad decisions cost my organization time and money," 70% of survey respondents either agreed or strongly agreed. Moreover, 82% of respondents either agreed or strongly agreed with the statement "Bad decisions harm my organization's reputation in the marketplace." Decision-making for customer engagement is business-critical—it could spell the difference between boom and doom!

Again, unlike rule-based systems that start breaking down as decisions become more complex, AI reasoning helps make decisions with a flexible approach, working with incomplete or ambiguous information and by comparing similar decision-making scenarios from the past. The output could be the final



Anand Subramaniam

Anand Subramaniam is senior VP of Worldwide Marketing for eGain. Prior to eGain, Subramaniam served in executive and marketing management roles in a range of organizations from cloud startups to companies such as Oracle, Autodesk, and Intel. He holds

an MBA from the University of California at Berkeley and an MSME from the University of Rhode Island.

decision or be one of the decisions in a larger customer engagement and contact center management process.

eGain AI clients decide:

◆ *An insurance giant decides how to precisely route customer queries to the right agent, with decision-making powered by eGain AI.*

◆ *A global media and legal services giant makes decisions on whether to take on malpractice lawsuits based on multiple factors like the amount of damages, win rates in the past for any given topic, tendencies of specific courts/judges.*

4. Comply with regulations

Enforcing compliance amidst ever-changing industry regulations, whether it is for content, access, or actual customer interactions makes customer engagement challenging for companies, particularly in highly-regulated industries. In fact, 70% of compliance professionals expect increased regulations, 59% expect increased personal liability, and 69% expect compliance staffing costs to rise, per Thomson Reuters. This is where AI can add immense value by guiding agents through the compliance maze.

eGain AI clients comply:

◆ *A marquee global bank uses eGain AI to guide advisors through a step-by-step dialog, and industry-compliant questions, to qualify and onboard small business clients.*

◆ *Another bank serves wholesale clients in multiple countries with a single eGain-powered AI system that incorporates the global parent's best practices, while considering local regulations in those countries.*

◆ *Utilities clients use AI guidance not only to resolve customer problems but also to help them avert safety hazards such as gas leaks, while complying with regulations on who asks and does what and when. This is not just mission-critical customer service—it is life-critical!*

Next steps

The best way to pilot AI is through a real-world trial in a production setting. That's exactly what eGain Try+Buy for AI offers, along with no-charge guidance to quick value and no obligation to buy! Contact us at info@eGain.com to learn more. ■

For more information on the companies who contributed to this white paper, visit their websites or contact them directly:



Accenture
Applied Intelligence
Contact: info.stc@accenture.com
Web: www.accenture.com/sca



Nuance
1 Wayside Road
Burlington, MA 01803
PH: 1.781.565.5000
Web: www.nuance.com



eGain Corporation
1252 Borregas Avenue
Sunnyvale, CA 94089, USA
PH: +1.408.636.4500,
+1.888.603.4246
Fax: +1.408.636.4400
Web: www.egain.com



BA Insight
7 Liberty Square, Floor 3
Boston, MA 02109
PH: +1.339.368.7234
Contact: sales@BAinsight.com
Web: www.bainsight.com/smarthub



Mindbreeze GmbH
Honauerstrasse 2
4020 Linz
PH: +43 732 60 61 62 - 620
Contact: office@mindbreeze.com
Web: <https://www.mindbreeze.com>



Sinequa
245 5th Avenue, Suite 1404
New York, NY 10016
PH: 1.646.560.8485
Contact: info@sinequa.com
Web: www.sinequa.com



Data Conversion Laboratory, Inc.
61-18 190th Street, Suite 205
Fresh Meadows, NY 11365
PH: +1 718.357.8700
Web: <https://www.dataconversion-laboratory.com/>



MC+A
332 South Michigan Ave., 9th FL
Chicago, IL 60604
PH: 312.585.6396
Contact: sales@mcplusa.com
Web: <https://www.mcplusa.com/>



Verint Systems Inc.
175 Broadhollow Road, Suite 100
Melville, NY 11747
PH: 1.800.4VERINT
(1.800.483.7468)
Contact: info@verint.com
Web: www.verint.com

Produced by:

KMWorld Magazine
Specialty Publishing Group

Stephen Faig
Group Sales Director
908.795.3702
sfaig@infotoday.com

LaShawn Fugate
Account Executive
859.361.0667
lashawn@infotoday.com

For information on participating in the next white paper in the "Best Practices" series, contact:
sfaig@infotoday.com | 908.795.3702 | lashawn@infotoday.com | 859.361.0667